REMARKS

Claims 1 – 20 remain pending in the application.

Objections to the Claims:

Applicant acknowledges the Examiner's remarks regarding the status of claims 1 and 10 as presented in the previous office action response. However, as the amendments appear to have been entered by the Examiner, Applicant is currently indicating the status of claims 1 and 10 as "Previously Presented" in accordance with 37 CFR 1.121(c).

35 U.S.C. § 103 Rejections:

Claims 1-5, 7-8, 10-14, and 16-20 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Chi, U.S. Patent 5,940,870, in view of Ang, U.S. Patent 6,678,799. Claim 9 was rejected under 35 U.S.C. § 103(a) as being unpatentable over Chi in view of Ang and in further view of Arimili, U.S. Patent Application Publication 2002/0112124. Applicant respectfully traverses these rejections.

The prior art references, taken singly or in combination with each other, fail to teach or suggest all of the elements of the independent claims. Chi teaches an address translation method for use in a system including a plurality of cluster nodes. The method includes, at a source node, receiving over a first network a communication with a first destination address having an index portion and an offset portion, wherein the index portion includes a partition number portion; providing an address mapping table which maps a plurality of indexes to a corresponding plurality of node ID's, each of the plurality of node ID's identifying a different one of the plurality of cluster nodes; using the index portion from the first destination address as an index into the address mapping table to identify a node ID, wherein the identified node ID identifies a destination node; appending the identified node ID to the first destination address to generate a second destination address; and using the second destination address to send information to a second network of the destination node.

Ang teaches a method and arrangement for cache management in a shared memory system. Each of a plurality of intercoupled processing nodes includes a higher-level cache and a lower-level cache having corresponding cache lines. At each node, update-state information is maintained in association with cache lines in the higher-level cache. The update-state information for a cache line tracks whether there is pending update that needs to be distributed from the node. In response to a write-back operation referencing an address cached at a node, the node generates difference data that specifies differences between data in a cache line for the address in the higher-level cache and data in a corresponding cache line in the lower-level cache. The difference data are then provided to one or more other nodes with cached versions of the cache line for the address.

In contrast, Applicant's independent claim 1 recites:

"a plurality of nodes, wherein each node includes an active device and a memory subsystem coupled to the active device;

wherein an active device in a node of the plurality of nodes is configured to generate a global address and translation information identifying a translation function, wherein the global address identifies a coherency unit;

wherein a memory subsystem included in the node is configured to select the translation function in response to the translation information and to perform the translation function on the global address to generate a physical address of the coherency unit within the memory subsystem;

wherein an additional memory subsystem included in an additional node of the plurality of nodes is configured to store the translation information identifying the translation function used in the node, wherein in response to a request for access to the coherency unit, the additional memory

subsystem is configured to send the translation information to the node" (Emphasis added).

Independent claim 10 recites a similar combination of features.

In the office action, the Examiner contends that Chi teaches an active device configured to generate a global address and translation information identifying a translation function and a node configured to perform the translation function in Fig. 8. The Examiner further contends that Chi teaches an additional node configured to store the translation information identifying the translation function in Fig. 8, #114 and #110, and col. 5, lines 27-46. Applicant respectfully disagrees with the Examiner's characterization, and submits that nowhere does Chi (including the portions cited by the Examiner) teach or suggest translation information identifying a translation **function**, selecting the translation **function** in response to the translation information and performing the selected translation **function**, or storing the translation information identifying the translation **function**, in combination with the other features as recited in the independent claims.

In col. 5, lines 27-46, Chi describes an address translation from the source node to the destination node:

"FIG. 8 shows the address translation from the source node to the destination node. The processor bus address 100 at the source node consists of two parts: an AMT Index 102 and an Offset 104. The AMT Index 102 is used to retrieve an entry from an Address Mapping Table (AMT) 106. Each table entry 108 contains a destination node ID for an access request on the processor bus. The number of entries in the AMT 106 is not necessarily equal to the number of the address partitions. For instance, in one embodiment, the AMT has 64 entries and the number of address partitions is 16. In that case, the AMT Index 102 is 6 bits while only the most significant 4 bits of these 6 bits indicate the partition number 110. In the AMT 106, the four consecutive entries corresponding to a given partition

number are filled with the same destination node. Note that the node ID from the AMT 106 is the physical node ID which is unique in the whole system. The partition number 110 is the logical node ID within each application, such that there may be more than one application running on the system."

Item 110 is described by Chi as a logical node ID within each application, and is also described as a partition number. Item 114, cited by the Examiner, is described as the node ID at col. 5, line 56. Nothing in the above citation teaches or suggests translation information identifying a translation function, much less performing a translation function in response to the translation information or storing the translation information identifying a translation function. Furthermore, neither Ang nor Arimili provide any teaching or suggestion that would remedy these deficiencies in Chi. Applicant thus submits that the prior art references, taken singly or in combination, fail to teach or suggest all of the elements of the independent claims, and thus respectfully requests removal of the 35 U.S.C. § 103(a) rejections.

Allowable Subject Matter:

Claims 6 and 15 were objected to as being dependent on a rejected base claim, but would be allowable if rewritten in independent form to include all of the limitations of the base claim and any intervening claims. Applicant appreciates the Examiner's consideration of these claims.

CONCLUSION

Applicants submit the application is in condition for allowance, and an early notice to that effect is requested.

If any fees are due, the Commissioner is authorized to charge said fees to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5181-95501/EAH.

Respectfully submitted,

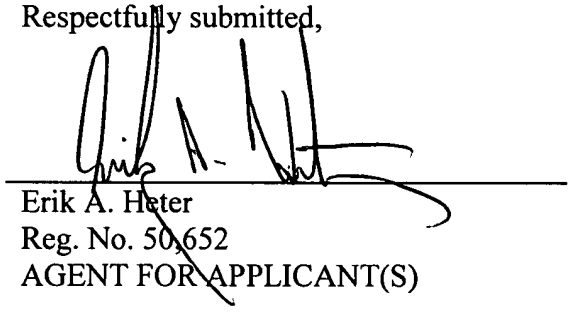

Erik A. Heter
Reg. No. 50,652
AGENT FOR APPLICANT(S)

Meyertons, Hood, Kivlin,
Kowert, & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
Phone: (512) 853-8800

Date: January 22, 2007